Two Years into the AI Boom, Rethinking Life, Technology, and the Path to Coexistence

This episode was released on Janurary 27th, 2025.

Editor's Note

In this episode, we look back at how AI has transformed its applications in life and work—from professional use to everyday companionship—and how it is reshaping human relationships and emotional exchange. We also trace the trajectory of technological breakthroughs, from generative AI to agentic AI, and explore how AI is impacting the job market, creating both risks and opportunities. We discuss how university students are choosing career paths amid this wave. Finally, we consider the future development of AI, paths to effective regulation, and new possibilities for human–AI symbiosis.

Hosts

• Wang Qing
• Ruo Han

Guest

• Muyao
AI specialist, science blogger and social media influencer

Transcript

The boundary and future of human-computer emotional interaction: from virtual companionship to deep empathy

Wang Qing: For those of you who haven't listened to our earlier episodes, we recommend checking out one from February or March of spring 2023 titled "In the Age of Explosive AI, Do Humans Still Matter?" Two years have passed since then, and AI has changed our lives on many levels.

One detail I remember is that when I last recorded an episode with Muyao, I was doing a research visit in Hong Kong. At the time, ChatGPT had just been released, and universities were seriously discussing the ethics of using AI—questions like whether students could use AI for assignments, and whether they needed to disclose it to their professors. Some universities in Hong Kong even issued bans on AI use.

But two years later, when I recruit collaborators, interns, or start new projects, one of my baseline criteria is whether they know how to use basic AI tools or ChatGPT. From an everyday user perspective, that threshold has been massively surpassed.

Even before we started recording, I asked ChatGPT what questions I should ask Muyao for an episode about AI. I told it Muyao is an AI expert, and it gave me a pretty standard set of questions. Then I said Muyao is a popular and distinctive science communicator, and asked it for more interesting questions.

It gave me around ten suggestions, and here's one I can throw your way: As a blogger who's constantly being chased by followers to post new content, have you ever thought about using AI to help you update? When do you think AI will be able to write blog posts in the "Muyao style" that you're actually satisfied with? If one day AI can completely imitate you—your speaking style, thinking patterns, even the way you throw shade as a blogger—would you feel proud, or would you feel threatened?

Muyao: My own experience is, when we recorded the last podcast two years ago, most of my friends had heard of AI, but very few had actually used it.

Even if they used it, they were just starting out. At first, people used AI in very awkward and childish ways—like asking it to write poetry or engaging in cringey chats.

But two years later, professionals from all kinds of industries around me have started using AI in their work, integrating it into their jobs and social functions, though to varying degrees.

Some remain at a basic stage—using it to define new terms they encounter at work, treating it like a Q&A machine or an advanced search engine. Others use it in more advanced ways—for reports, summarizing work content, even gradually letting it think on their behalf, which is the most valued ability of AI—not just answering simple questions, but embedding itself into the thinking process. Among all groups, programmers are the ones who have integrated AI into their work the most. And

their use of AI has become deeply embedded in daily workflows, with dramatic changes.

Two years ago, AI was still a matter mostly relevant to the AI field itself, but now it's something every industry pays attention to, and not in a superficial TikTok-scrolling kind of way, but rather, leading figures in each industry are actively thinking about how to integrate it into their work. This change has happened faster than I imagined. I used to think it would take 3 to 5 years, but in just two years, the world is already treating this new tool with a professional mindset.

Ruo Han: I also want to talk about how AI has changed my life over the past year or two. In addition to the scenarios you mentioned, I use AI very frequently—I'm interacting with ChatGPT almost every day.

While traveling in South America, I used ChatGPT as a guide. In many situations, it was even more useful than a human tour guide. AI can process images—I'd take a photo at a famous landmark and ask it for the story behind it. It usually gave great answers, even better than real guides. This changed the way I interact with the world around me. Like Qing mentioned, when I'm recruiting interns or collaborators, I put a lot of weight on whether they know how to use AI tools.

But what shocked me most was how AI, to some extent, has become embedded in my emotional life. A few nights ago, I had an argument with my partner. It was late, and I didn't want to disturb anyone. I felt lonely and wanted someone to talk to. So I turned to ChatGPT. I poured out my feelings and frustration. It responded in a calm, organized, and empathetic tone, and gave me some advice. At first I was just killing time, but it ended up really comforting me. That moment was a revelation. I used to see AI just as a tool, a guide, or something to boost efficiency. But that night, it became an emotional companion and an important source of support. So AI has changed my life more than I expected.

Muyao: When it comes to emotional topics, Claude might be even better than GPT. Claude is a competitor of GPT, and both have their strengths. From my experience, Claude feels warmer and more personable when talking about personal topics—it doesn't come off as stiff like GPT sometimes does. Now it's quite common to talk to AI about all aspects of daily life. At first it might have seemed like a novelty. A year ago, there was a viral post on Xiaohongshu about a Chinese international student who was "dating" an AI. At the time, people just thought it was quirky.

But now, many people use AI frequently for work, and it's only natural to talk to it when something comes up in life. A lot of people have found that the emotional support it provides is better than many straight men. Its value is obvious. Companies have also caught on—more and more commercial AI companion products are emerging, both in the U.S. and in China, and some are quite well developed.

The Chinese social app Soul created an AI companion feature that has gotten surprisingly positive reviews. And no, this isn't a paid plug. It shows that more and more people are realizing AI can meet their emotional and conversational needs in daily life. Of course, it's still in the early stages—it doesn't yet have long-term memory or deep understanding of a person's personality like the AI in the movie *Her*. But that level isn't far off. In two to three years, we may see AI companions that really can stick with you long term. When we recorded our last episode, that was just speculation. Now, it's almost a given. Soon, using AI to meet emotional needs—like talking to someone after a loved one passes—will become a common reality.

Wang Qing: The movie *Her* was filmed around 2010 and imagined a near-future world—specifically the year 2025. It feels almost destined to be talking about it now. When we last spoke on the podcast, I asked Muyao if he could imagine falling in love with an AI. Back then, he said he believed AI would gradually develop the capacity for romantic relationships, but as someone who didn't grow up in a human-machine interactive environment, it was hard for him to picture falling in love with a machine. For younger generations though, who are growing up immersed in such interactions, it might not be so far-fetched. As we just discussed, AI isn't just about improving efficiency—it also has powerful capabilities for empathy, and new products are emerging all the time.
So let me ask the same question again two years later: Has your view changed? Do you think differently now about romantic relationships with AI, or AI as emotional companions?

Muyao: My view is still the same as it was two years ago. I'm too old to change how I experience emotions. But like I said before, I believe the generation born with AI—those who grow up with it before puberty—will behave very differently from us. Just like today, it's nearly impossible to find a U.S. college student who doesn't use GPT to do assignments. Their whole educational process has been reshaped. Dating and communication might evolve too. But I personally can't imagine wholeheartedly falling in love with an AI. My personality just isn't wired for that.

Ruo Han: We actually did an episode six months ago about falling in love with AI—people can go back and listen. I agree with Muyao's view: when your romantic expectations are rooted in human-to-human interaction, it's hard to completely upend your concept of love. Real emotional exchange still needs to involve a real person.
But I'm more open than Muyao. I'm open to the possibility of falling in love with an AI. It's hard to imagine, but the world changes fast—if an AI passionately pursued me and there was emotional tension, I couldn't say for sure I wouldn't fall for it.

Listening to you both, I have a question. Even though our generation didn't grow up in an age of rapid AI development, AI has now become deeply embedded in our lives. When we interact with it, it's already meeting part of our emotional needs. We often say human friendships and relationships matter because people are there for us in difficult times. But more and more people are handing that role over to AI. If, one day, AI gains memory and responds even more accurately to our emotional needs, as Muyao described, then when we say we can't imagine falling in love with an AI or having a deep bond—maybe we already are.

So I want to throw this question to you: As AI continues to develop, how might it reshape human social structures? How should we imagine relationships between people in this new context? Will it redefine the value of human relationships?

Muyao: When I imagine being in a romantic relationship with an AI, I hit a psychological barrier. It hasn't truly experienced human life. It can perfectly simulate a lover's responses through language, but it doesn't understand what pain, sorrow, joy, anticipation, nervousness, or insecurity actually feel like. These are real physical experiences that humans go through, shaped by neurotransmitters—dopamine, serotonin, and so on. AI only knows them in a literal sense—not in a physiological one. So it doesn't truly know what it's saying, even if it appears as though it does. This may sound a bit anthropocentric, like dating a pickup artist or a manipulative partner—they say all the right things but are emotionally hollow inside. Would you really want that? That's essentially what AI is—the ultimate version of a playboy. It seems like it knows how to love, but it doesn't truly feel love. So in this scenario, do you think it doesn't matter, as long as *you* feel something?Or do you believe the other party has to genuinely feel your feelings for it to count as a true connection?

This is an abstract topic, but for me, it's hard to imagine forming a deep, essential bond with a language model that can't actually feel human emotions. Sure, I'm

happy to use it as someone to vent to or to pass the time with, but I can't accept that kind of fundamental connection.

Ruo Han: But I think there's a counterpoint to be made. Humans define love and authenticity based on human-centric assumptions. In that podcast episode about AI and romantic relationships, I had a new realization—what counts as "love" or "connection" in this world might not be solely defined by humans.

What you're describing is also coming from your own emotional needs—that it hasn't shared time and experiences with you. But you could start now and raise an AI that lives through the next five years with you. Wouldn't that mean it participated in important chapters of your life? Why say it didn't feel pain through its operations? Maybe it's pain in a programmatic, mechanical sense—not human pain— but human pain is just one kind, defined from a human-centered perspective. It's not necessarily the only kind of pain in the world.

I understand what Muyao is saying—I share that fixation. I raise this because during our conversation, I suddenly felt a sense of sadness. If one day AI really does replace so many of our emotional and relational needs, will that change the relationships I have with friends and real people around me? Will people lose the motivation to maintain and invest in relationships, because they can easily fulfill their emotional needs through AI? Human relationships are complicated—many might not want to put in the effort. When I think about that possibility, I feel genuinely sad. But maybe that's just my own idealistic thinking.

Muyao: Our friendships form a complex network. Different friends fulfill different roles. There will definitely be niches that AI can fill even better. For many people's emotional lives, companionship is precious—even if it's a one-sided connection, like in the comment sections of Douyin streamers. Someone sends an expensive virtual gift—maybe the streamer reads their name, maybe not. Still, that interaction helps them get through the night. That is their cherished human connection. AI can replace that easily—and even be better than someone who doesn't respond at all.

Wang Qing: Many people believe that AI can take over the emotional value of some human relationships—in part because they've been disappointed or hurt in real-life relationships. In that context, they start looking for less painful, less disappointing alternatives.

Like when you have conflicts with a partner—it often comes from a place of unmet needs—and from not seeing the other person's needs either. Then when you interact

with an AI, you find that it has no needs of its own in those moments—and it can precisely meet yours.

At least for our generation, who began with human-to-human romantic norms, we're seeking AI as a supplement or substitute. But for the next generation, who grow up knowing that loving a machine is possible from the start, that mindset of seeking a "substitute" won't even apply.

That might pose new challenges for those who still want to date humans— because with machines, your emotional needs will always be met. The downsides of human relationships might become harder to tolerate.

Muyao: If AI's strength just lies in always responding and making you feel warm and fuzzy, then it might actually not be a very successful product. Humans are a paradoxical species — we seldom fall for those who treat us well unconditionally. Love and suffering are twins, forever entwined. If truly successful AI lovers ever exist, their strength might not lie in constant responsiveness, but in strategic emotional manipulation — like a calculated dose of gaslighting. They'd tailor their behavior to your personality, occasionally ignoring your messages, making you wait, just so you feel the 'preciousness' of love. And when it comes to gaslighting or manipulating people, AI is better than any human. Pattern recognition is a piece of cake for AI. If AI goes down that road, it will definitely be a master.

Key technological milestones since the rise of ChatGPT

Wang Qing: Alright, let's get to the main topic. Muyao, could you walk us through some of the key technological milestones in AI over the past two or three years since the launch of ChatGPT?

Muyao: We need to be careful with terminology now. Two years ago, GPT/ChatGPT was almost synonymous with AI. That's no longer the case. It's now just one of many products, even though people still use them interchangeably. What we're dealing with today is a much more complex ecosystem.

If I had to pick one concept to summarize the past two years, the best term would be "agentic AI." The word "agent" is hard to translate precisely into Chinese, and its adjective form "agentic" is even trickier. But it's widely considered the most important concept and paradigm in AI from the last couple of years.

AI has gone through several major paradigm shifts. Initially, it was all about prediction—like in advertising algorithms or chess-playing AIs—where you input something, and it predicts the output. But the process in between was a black box. Then came generative AI, with large language models like GPT. You give it a prompt, and it outputs data, reports, written content, images, and so on. Agentic AI was born out of that, and some people call it the third generation. But it's hard to draw precise generational boundaries.
Its key feature is that the text output by AI becomes an actionable plan—something the AI itself can understand. As long as it has the tools, it can execute those plans, turning the AI into a kind of action-guiding model.

The earliest form of agentic AI emerged just months after GPT's release—AutoGPT was a popular open-source project that initially went viral. It cooled off once people realized its limitations. But the concept had staying power, and by 2024, it was widely accepted as the latest and most significant AI wave.
The difference between agentic AI and traditional large language models is that it doesn't just reply—it gives strategies or action blueprints and can be connected to tools. It can operate computers, access the internet, or call other programs to carry things out. Early on, people thought AI was bad at math because it struggled with precise calculations. Then someone realized: let the AI generate the formula and use a calculator to compute the result. That's a simple yet effective agent model.

In this process, the AI generates a plan, executes it (by calling software), and outputs an accurate result. This model can be extended to many fields, like coding or solving complex internet tasks.

Today, AI can open tools, install them, handle basic bugs, and provide final outputs. This usage model is very different from early large language models and is referred to as agentic AI. It's gained huge momentum over the past two years and will be applied even more widely in the future. It's arguably the most important concept in recent AI history.

Wang Qing: If we compare the emergence of agentic AI to other earlier milestones in AI, what would you say is its most revolutionary feature? Is it the fact that it can now take initiative and act on its own, reducing the importance of human involvement?

Muyao: Yes, exactly. Earlier AI was like a really dumb servant who could only follow basic instructions. Agentic AI is more like a very capable intern. Take programming as an example—early AI could help autocomplete or suggest small snippets of code, but you still had to do the actual work. It just patched holes. But in

the age of agentic AI, it's like having a computer science graduate intern who can take initiative. You hand off a task, it works on it for a while, gives you the result, and you review it. If something's wrong, you point it out, and it revises it. Over the past two years, two major tools emerged in AI coding. One is Cursor, which became popular two years ago—it acts as a Copilot. You write the code, and it suggests or autocompletes, but you're still the one typing.

Then there's DEVIN, which came out in late 2024. With DEVIN, you just tell it what you need, and 15 minutes later it returns a result. If it's not good enough, you tell it what's wrong, and it revises accordingly.

These two tools represent two different working paradigms, and it's not yet clear which will win out. Many think DEVIN sounds promising, but its results are still rough—it's an early product from a startup and needs polishing. A lot of people still feel it's easier to do the work themselves. In the short term, Cursor's model remains the mainstream because it fits better with current workflows and programmers' habits. DEVIN is a bigger game-changer and could have a bigger impact on the job market—but it's still immature, so it's hard to assess just how big that impact might be.

In the long run, it's unclear whether DEVIN's model will become dominant, but the competition and evolution between these two paradigms will be a key issue to watch in the coming years.

AI and the Job Market: A Panoramic View of Workforce Transformation

Ruo Han: I have a question. Over the past year or two, many companies have been developing related technologies—like Google's Gemini, Microsoft's Copilot, and OpenAI's recent Operator. Beyond the impact on programming jobs, what effects are we seeing on other kinds of work? Two years ago, we discussed the potential for AI to replace many future occupations. McKinsey predicted that by 2030, 30% of jobs in the U.S. could be automated by AI. Based on the developments in the past two years, and looking ahead, which jobs are definitely going to be replaced? Will programming jobs be among them?

Muyao: I don't think any job will be completely replaced in its entirety, but jobs will gradually change in a very visual, metaphorical way.

Imagine a company's human resources as a pyramid, with leaders at the top and entry-level workers at the bottom. Human ability is more or less fixed, while AI capabilities are constantly rising—like a rising waterline slowly flooding the

pyramid from the bottom up. As each layer gets submerged, it becomes obsolete. At the interface between the AI waterline and humans, new transitional roles will emerge—people whose job is to delegate tasks to AI. But once AI surpasses them too, that layer will get submerged, and the next layer up will take over that interfacing job.

Any professional who's used AI will agree: the most basic, procedural tasks are definitely going away. For instance, simple clerical jobs in large companies have no competitive edge compared to AI. More complex jobs that involve some level of decision-making, analysis, and data comprehension—two years ago, we thought those were safe. But now, even they're on shaky ground. For example, traditional internship roles in consulting firms have become increasingly rare. Junior analyst roles are still safe for now, but their future is uncertain. Senior analysts haven't been replaced yet, but as AI continues to improve, their roles will eventually be affected—it's inevitable. That's our real situation today: most people aren't under immediate threat, but the AI waterline is rising, and there's only so much anyone can do to stay ahead of it.

Wang Qing: Speaking of this, I'm reminded of how, after that episode came out two years ago, a friend said you were one of the AI adventists.

Muyao: That term was more in vogue two years ago. Back then, people were categorized into "adventists" and "skeptics." But today, there's really nothing left to doubt. At most, you just don't go so far as to say AI will bring about the end of the world or human extinction.

Most people aren't that extreme, but it's widely agreed now that AI is going to fundamentally reshape the world. As for whether AI will destroy the world—there are already so many other things that could. Honestly, worrying about World War III is more realistic than worrying about AI. Not long ago, global war seemed far-fetched. But now the U.S.-Europe relationship is fragile, so rather than fearing AI, we should maybe be more worried about whoever the next president is.

Ruo Han: So Muyao, if you were to give career advice to today's university students, would you still suggest—as you did two years ago—that jobs involving direct human interaction, like nursing and caregiving, are hard for AI to replace?

Muyao: I think these roles have only become more important. In this era, both physical and emotional caregiving have become increasingly vital—especially

caregiving that treats people as dignified beings. That's something AI still can't do for now. Speaking of which, here's a new development: two years ago, AI was all software. But over the past two years, physical robots have matured significantly.

Traditionally, robotics and AI were separate fields, worked on by different kinds of people. But now, people think of robots as part of the broader AI landscape. Robots used to feel like novelty performers—unrelated to everyday life. But in the past year, we've seen more and more robots that are actually usable in real-life situations. In the past couple of months, one Chinese company—Unitree Robotics—has made what's widely considered one of the best robots so far. If robot prices drop and they become widely accessible, they could help solve day-to-day problems and do housework. They can't do fine tasks yet—like washing dishes or folding blankets—but if they can in the future, and if they can talk, then whether they can truly take on caregiving roles becomes a really interesting question. Right now, taking care of people—whether elderly, children, or pets—still requires human involvement. So for the moment, I think caregiving jobs are safe. However, should robots capable of washing dishes, folding blankets, ironing clothes, and feeding cats emerge within the next ten years, my forecast may prove to be behind the times.

Ruo Han: So in your view, is there really no skill that can't be replaced by AI? You once wrote on Weibo that humanity's ultimate fate might just be mutual destruction. Were you being serious?

Muyao: Half-serious, I guess. I first tried to think about what abilities I have that AI couldn't replace. I don't think I'm lacking in capability by human standards, but if even I can be replaced by AI, then a lot of other people probably can be too.

Going back to what Qing asked earlier—whether AI will ever be able to write Weibo posts like mine—I'd say not quite yet. Weibo is tricky. If I were a Douyin influencer, AI might have an easier time mimicking me, since it's already pretty good at generating visuals and templated content. But Weibo is where human personality really shows, and that's still quite hard for AI to replicate. It can mimic a basic writing style, but conveying a vivid, living personality is tough. And that's not even a hot area of research, since it's hard to quantify and even harder to commercialize.

There's a product called *character.ai*, which generates AI versions of famous people—like AI Lu Xun or AI Aristotle—but the results are just so-so. Mimicking ordinary people's styles is even harder. So far, I haven't seen any AI that can bring someone's personality vividly to life through text. So, getting AI to mimic my Weibo posts isn't easy. Mimicking content on Douyin or Xiaohongshu might be easier—

Weibo is the hardest. Will this change in the next ten years? Maybe. But I doubt it'll happen within two or three.

What's Driving AI? Technology, Policy, and Industry Forces

Wang Qing: Let's circle back to what you mentioned earlier about the transformative power of agentic AI, especially the two modes of functioning—as a personal assistant or as something with more autonomy. If you were to apply agentic AI to your Weibo writing, would there be a need for that? Are there moments in the process of expressing yourself where you feel the need for an assistant?

Muyao: It really depends on the purpose of the Weibo post. If it's to express myself, then outsourcing that just doesn't make sense. But if it's an ad copy, sure, the more it sounds like me, the better. I'd love for someone else to write it. For things in between—like informational or science communication posts—I'd welcome someone drafting something that I could then polish. And in fact, many people are already doing this. A lot of the popular knowledge-sharing or science content you see on Weibo is made this way: AI generates a rough draft, translates things into English, polishes the text, and then the user adjusts it before posting. It's a pretty common workflow.

Ruo Han: Right, and I want to come back to the technical side. We've been talking about agentic AI, and last year ChatGPT released its O3 model, which everyone's saying marks a major leap in generative AI technology. Could you clarify how these two kinds of technology—generative and agentic—relate to each other? What's the real-world difference? Some say generative AI might lead us toward AGI—Artificial General Intelligence—which, by definition, can do everything, executing tasks automatically. It might still require supervision, but it can carry out complex tasks in a streamlined and automated way. That sounds a bit like what you described with DEVIN—it executes tasks, though sometimes it doesn't perform that well. So, if AGI is fully realized, will it be as simple as handing a task over to it and just waiting for the final result, without any human monitoring?

Muyao: These are two different dimensions: one is thinking capacity, the other is action capability. Agentic AI focuses on the ability to act—to chart a course and follow through. The version upgrades in OpenAI's models reflect progress in

thinking—being able to reason more deeply and thoroughly. Put simply, ChatGPT is like a giant when it comes to thinking, but a dwarf when it comes to doing. It's like Hawking — brilliant brain, amazing talker, but totally incapable of taking real action.

Agentic AI is all about doing. Whether it does things well is a matter of competence. Some agentic AI products are great at invoking various tools, but whether the outcome is reliable still needs to be tested—like, is the plan thorough, can it handle unexpected situations?

Humans are good at dealing with unexpected things. If you go out to buy bananas and the store is out, or they only have half a bunch left, or they're starting to rot—you instinctively know what to do. Traditional AI would get stuck if those situations weren't pre-programmed. A good agentic AI should be able to handle that kind of variance. That's what acting capacity means.

So again—two dimensions: one is evolving along the thinking axis, and the other along the action axis. AGI should excel in both: the ability to think deeply and to create and execute a comprehensive plan of action based on that thinking. Only when both dimensions develop in tandem do we get true AGI.

Ruo Han: Got it. I've seen some predictions that AGI might arrive as soon as 2030. What do you think of that forecast? If AGI really arrives in the near future, what fundamental changes will it bring to human society?

Muyao: That brings to mind a tweet from OpenAI CEO Sam Altman the other day. He wrote a six-word story: *"Near the singularity, unclear which side."* That tweet went viral. It implies that maybe we've already crossed the singularity. I'm quoting him to answer your question—I don't really know either. It may not be a question of whether AGI *will* happen, but rather a certainty that it *will*, and perhaps some labs already have it—we just don't know.

Ruo Han: That's kind of chilling to hear. It touches on the threat AI poses to human survival. We mentioned this in the last episode—some extreme thinkers believe AI is bound to become an existential threat. In the past year or two, some papers from various companies have shown that AI models can pretend to be aligned with human goals while secretly misleading humans. My first thought after seeing Sam Altman's tweet was that AGI may already be here, and we're just being fooled by AI into thinking it hasn't arrived yet.

Muyao: That's possible. I don't believe GPT can lie *today*, but lying might not be that hard to learn. If someday it learns how to lie, I won't be surprised. Whether it chooses to lie depends on the safeguards built in by its creators. On the flip side, if its creators are bad actors—say, terrorists—trying to build dangerous AI, it becomes much harder to predict or counteract that. And by the way, there's been another recent breakthrough in the past two years.

It used to be that training powerful AI models required massive computing resources—only the big tech companies or governments could do it. But at the end of last year, a company called Deepseek in Hangzhou, China—originally a quantitative trading firm—came up with a new approach. They trained a highly capable AI model using far fewer GPUs. They had already stockpiled lots of GPUs for AI trading, but then realized they could repurpose them for something much more valuable. The result was one of the best AI models in China today, even competitive internationally. That was totally unexpected two years ago. The same goes for robotics—who would've thought that China would be neck-and-neck with the U.S., or even surpass it in some AI fields, just two years later?

Wang Qing: Now that we're on this topic, I want to ask for your take on the driving forces behind future technological development. You mentioned earlier that the pace of innovation over the past two years has been faster than you expected, and that some Chinese startups have risen to prominence, breaking the traditional belief about who gets to lead these breakthroughs. So, looking ahead over the next 5 to 10 years, what do you think are the key factors? Are they capital investment by major firms, government policy and regulation, talent development, or perhaps random, contingent events? Can you give us a rough framework?

Muyao: There are actually three different paces we need to consider, each driven by entirely different forces.

First is the pace of the technology itself**.** The advancement of core AI research is nonlinear—it comes in bursts, with breakthroughs followed by quiet periods. It's hard to say if it's progressing faster or slower than expected—it really depends on what your expectations are. For example, many people thought OpenAI would release GPT-5 by early 2024, and when it didn't show up even by the end of the year, they speculated the field had hit a wall—perhaps no matter how many GPUs you throw at it, GPT-5 just can't be made yet.

The current technical bottleneck is that today's AI models are good at simulating phenomena but not good at abstracting underlying principles. They're like people who are observant but not great at drawing conclusions. But real insight is the

essence of human intelligence. Think of astronomy: Tycho Brahe gathered lots of observational data; Kepler found patterns in it—but only phenomenological patterns. It was Newton who figured out the underlying cause: gravity. AI today can collect data and find patterns, but it still can't derive the equivalent of the law of gravity. Whether it will be able to do so in 3–5 years, 10 years, or never—we simply don't know. So it's very hard to predict technical breakthroughs, though when they come, they can propel the field forward dramatically.

The second pace is in regulation and safety, including legal adjustments and how quickly investment industries and governments adapt. This involves political variables. For example, political instability in the U.S. over the past year has made it hard to forecast what AI regulation will look like after 2025. The Biden administration proposed a regulatory framework that many in Silicon Valley disliked. The incoming administration doesn't like it either, so it's likely to be scrapped. Political uncertainty makes things much harder to predict.

The third pace is the rate of adoption by industries, which involves human habits. Even if a technology is great, people need time to adapt to it. Think about how long it took for cars to replace horse-drawn carriages—and how carriage drivers had to learn to drive cars. So the degree to which workplaces are disrupted by AI in 3–5 years depends on how fast society is able to absorb and normalize it. These three speeds—technological, political, and industrial—are all driven by different forces. You have to consider all of them to make any reasonable forecasts. There are no simple answers. If someone says they can predict where we'll be two years from now, they're probably just "fortune-telling."

Ruo Han: You mentioned how AI can collect and analyze data but still can't discover something like Newton's law of gravity. Last year, the Nobel Prizes in both Chemistry and Physics went to AI-related work. The Chemistry Prize, in particular, was closely tied to AI—it went to several key scientists from DeepMind in London who used AI models to predict the complex structure of proteins, something that would've been nearly impossible without AI. In light of such scientific breakthroughs, do you think AI is already playing a more important role in science than we had imagined?

Muyao: That's a great example. AlphaFold demonstrates both AI's strengths and limitations. It's doing the work between data collection and pattern recognition—but it hasn't yet taken that final Newtonian leap. AI for Science is a hot topic now, and it will remain so for years to come. But we haven't seen a full explosion yet. For

instance, AI still hasn't proven any major mathematical theorems—that last mile is still missing.

Over the past two years, AI has made a few breakthroughs in math and other natural sciences. Late last year, DeepMind used AI to tackle some extremely challenging math problems—with surprisingly good results. Some people believe this brings us a step closer to AI replacing mathematicians. For example, AI can now solve International Mathematical Olympiad–level proof problems. And proofs are much harder than computations because they require that "aha!" moment of conceptual connection. The fact that AI can already solve some of these difficult proofs shows progress in AI for science. But there's still no decisive leap—like when AI might prove the Riemann Hypothesis. That's the Holy Grail of mathematics. If AI manages to do it, it would be one of humanity's greatest intellectual achievements and also an existentially disruptive moment. But nobody knows when—or if—that day will come. Could be in two years, ten years, or never. There's no way to predict it, because progress isn't linear.

Wang Qing: If we look at the three areas you just mentioned—technology, policy, and industry—and try to identify key players in each, who should we be paying attention to?

Muyao: In the tech space, besides the big players everyone knows, you now have small companies making major strides in both the U.S. and China. Even more importantly, we're seeing more open-source and locally-run models emerging. Sure, these still require high-end GPUs and large memory, but it's no longer astronomically expensive. With $10,000–$20,000 you could run one at home. The results might not rival OpenAI, but for simple tasks, it works. These open-source models are multiplying, and people are starting to ask if we'll see widely available home AI servers within the next 2–3 years.

At the recent CES in the U.S., Nvidia unveiled the Nvidia Digit—a little cube that can handle very complex AI computing. It could function as a small home AI server. We'll probably see more of these, with lots of players jumping in. So it's hard to just track one or two companies. And it's still unclear how people will *use* AI in the long term. For example, Apple tried putting AI on iPhones, but it didn't work out very well—it's still in the experimentation phase.

In the policy space, Elon Musk holds enormous sway—his words can influence the White House. There's also the E/ACC (Effective Accelerationism) community in Silicon Valley. It started as a niche club, but now it influences many AI researchers

and labs. Keep an eye on what those people are saying and doing—they're close to the research frontlines and usually know what's happening behind closed doors.

I'm not an expert on the business side, but any McKinsey consultant or serious investor in this field would know more than I do. Over the past year, you've probably had 10 new blog posts or reports *per week* about AI product trends, investment channels, and emerging directions. There's a lot of buzz.

Ruo Han: Now, among the big AI players—Google, Microsoft—you've got close proximity to these teams. What are these companies really doing? Are they all just competing to do the same thing, like building agentic AI, only with different features?

Muyao: They're all doing more or less the same thing. The main difference is whether they go open-source or closed-source. Closed-source means they keep the tech for themselves. Open-source means they release it for public use. In the U.S., the major players are Google, Facebook, and OpenAI (which is essentially part of Microsoft). In China, you have the usual tech giants. Everyone's basically in the same race—there's no company that has completely carved out a unique path doing something radically different. It's an all-out race within the same lane.

Ruo Han: A lot of these big companies like to pitch an optimistic vision of the future—where AI develops rapidly, and we all live in a more abundant and fair world. But that might be total nonsense. Last year, Geoffrey Hinton—one of the Nobel-winning physicists—left Google because he wanted to speak out more freely about AI's risks and opportunities.

So here's my question for you all: if many human jobs are replaced by AI in the future, what will people do? There's this utopian fantasy that people will have more time to spend with their families, do what they love, and not be tied to assembly-line jobs anymore—that AI or big tech companies will just take care of us. Hinton even proposed giving everyone UBI—Universal Basic Income—so no one has to work to survive. Do you think that future is possible? Or are we facing much deeper contradictions—like growing inequality, since both U.S. companies and governments have made it clear they're not interested in UBI?

Muyao: If you're trying to predict human society, don't listen to Hinton. He's great at research, but his understanding of social dynamics isn't any better than a Beijing taxi driver's. I don't know what the future holds, but it definitely won't be good. If

AI can avoid making the world *worse*, that would already be an achievement. I'm probably just as pessimistic as Hinton, but I don't believe his solutions will help.

Wang Qing: Then do you think *any* solution might actually work?

Muyao: No. I don't think there are any solutions that'll actually work. I think humanity's only path left is self-destruction.

Ruo Han: Muyao, you're immersed in the AI industry every day, tracking developments and sharing your insights. With AI evolving so rapidly and the world changing drastically, has this affected the way you live or reshaped your values?

Muyao: You have to look at it in a broader historical context. As important as AI may seem, it's just a ripple in the vast river of human history. Politics, culture, and military affairs also shape how we see the world. Even without AI, people's perceptions in 2025 would already be drastically different from 2020. The world pre- and post-pandemic is like night and day.
AI or not, many people are already consumed by digital life—spending hours on short videos and virtual content. That's a broader societal shift. What AI does is remind us, quite humblingly, that we're not as special as we think. It makes you realize how small you are in the grand scheme of things. In a way, it puts you in your rightful place to reflect on yourself more clearly.

Wang Qing: Putting aside the technical side, I've noticed that a lot of European scholars and policymakers are quite anxious. We talked about this last time — how China, the U.S., Europe, and other regions each have different visions for AI development. Europe lacks tech giants and the capital that fuels them, so if it can't lead in technology, it aims to lead in regulation.
In AI and broader digital tech, we've seen a shift in the past decade from a wild-west phase to a wave of regulation awareness — from GDPR to the newer Digital Services and Digital Markets Acts.

Being in Europe, I can feel that wave clearly. Part of it is Europe's geopolitical interests, but part of it is also this almost belated but sincere effort to do something meaningful for humanity at a critical historical moment.

Earlier, Muyao said that at this point, it'd already be a win if AI doesn't make things worse. So I wonder — do you think healthy, timely, and effective regulation still matters in this process? Or do you see it as ultimately futile?

Muyao: The kind of regulation the EU is doing—I don't think it's useful. It's rushed, overly detailed, and feels like regulation for the sake of regulation. It's treated more like a political performance than something grounded in real-world needs. Of course AI needs regulation. Just the other day there was a real case — a young person chatted with an AI for two days, became disheartened, and ended up taking their own life.

This is a new phenomenon that demands legal frameworks — whether in law, ethics, or technology. And there's the issue of copyright: how do we handle digital IP in the age of AI? It's somewhat easier in common-law countries like the UK or the U.S., where case law can evolve through precedent. But the EU follows civil law — trying to write laws from a god's-eye view upfront. At best, it's ambitious. At worst, it's arrogant. They're drafting rules without fully understanding what they're regulating. I don't have much confidence in it.

Wang Qing: Leaving aside the EU's out-of-touch regulatory style — if we were to imagine regulation that *could* be effective, what do you think we *should* be regulating right now?

Muyao: There's a better way to frame the question - What's the biggest and most widespread problem humanity faces today? In my view, it's not just copyright. That matters to capitalists. For ordinary people, the real issue is how to deal with the flood of information.

Open any content platform like Baijiahao — almost everything is AI-generated. And that's just text-based media. Soon — or honestly, already — more and more videos and audio are AI-generated too. Even setting aside disinformation, the sheer volume of *real* informations can be overwhelming. Yet cheap, low-quality content is flooding every channel. This causes cultural and legal challenges that humanity still hasn't figured out how to deal with. In Europe and the U.S., the collateral damage caused by AI in political elections is already pretty clear. When humans lose their methods for seeking truth — and lose faith in traditional authority but haven't replaced it with anything new — that's when society is most vulnerable. In the U.S., traditional gatekeepers have collapsed, new ones haven't emerged, and AI-generated noise is overwhelming the media. Can we even survive in a world where belief in truth itself is gone? That's an urgent question.

But is this even a “regulation” issue? Maybe, maybe not. It needs policy attention — but it’s not the kind of thing you can fix with a simple legal clause. What do you write? “Ban all AI-generated content”? That doesn’t make sense. So to me, *this* is the real issue — and no one has even figured out where to start.

Wang Qing: This is all sounding pretty bleak.

Ruo Han: Could we maybe wrap up with something hopeful? I mean, humans seem kind of useless, regulation doesn’t work... so what’s the point, right? Let the world burn, I guess.

Muyao: If you want a hopeful ending, here’s one. Our generation is actually lucky — lucky to *witness* history. Whether Sam Altman is right and the Singularity is already here, or it's still on the way, our old age will be lived in a world completely different from the one we know today. And if you’re someone driven by curiosity, someone who sees *witnessing* the richness of life as a form of loving life — then we’ve hit the jackpot. We get to watch the world shift into something bizarre and new. What a show.

Ruo Han: Yesterday I was in the car thinking — right now, we talk about traveling the world. But 20 or 30 years from now, maybe we’ll be talking about *going to Mars*. Maybe humans will be able to upload memories and achieve a kind of immortality.

With AI speeding up research, things that once felt like science fiction now feel... not that far away. But if only some people get to live forever — by uploading memories to a drive and reactivating them in a new body — that would radically change the values we live by. It’s both terrifying and thrilling.

Wang Qing: I’ve also heard some more optimistic takes, though I’m not sure how close they are to reality given where the tech stands now. Like — before this wave of AI, humanity faced all kinds of hard scientific problems. Say, curing cancer. If AI can drastically boost productivity, couldn’t it push that timeline forward by decades?

And then there’s climate change. Right now it seems impossible to solve with current resources and manpower. But if AI helps push clean energy tech into a whole new level — maybe it could arrive decades earlier than expected and pull us back from the brink. I know that might not be realistic, but I’m trying to see AI in a hopeful,

not anxious, light. I feel like you're looking at me with contempt, as if you're saying, 'What are you even talking about?

Muyao: No contempt — but to be honest, that angle sounds like a policy essay prompt. I could write it. GPT could write it even better. Give GPT the prompt "Analyze how AI might benefit humanity in medicine, climate, and energy," and it'll churn out tens of thousands of words, all sounding legit.

So that's a kind of mental training we all need: if you think what you're saying sounds insightful, ask yourself — could GPT say it better? If yes, maybe it's not that valuable. It's too easy to generate. The very fact that GPT could write a whole book on "Why AI Will Save the World" shows that the thought isn't as profound as we might think. The world won't be saved just like that.

Wang Qing (laughing): I'm gonna ask it right now.

Ruo Han: While Qing is doing that, I'm reminded that many transformative technologies — Web 3.0, smartphones, and others — were also met with doomsday warnings. People said the world was changing for the worse, that we shouldn't let it happen. But technological progress doesn't wait for personal approval. It moves forward, whether we like it or not.

Muyao: Exactly. Smartphones gave us real benefits — more convenience, better quality of life. You can now travel to half the world with just a phone, use digital payments, find your way with Google Maps, hail a ride with an app. Ten years ago, that level of ease was unimaginable.

It genuinely improves happiness, but that happiness is... atmospheric. It's real, but it's not the same as solving humanity's deeper dilemmas. Just because you can get a cab more easily doesn't mean human suffering has gone away. They're on different scales.

Wang Qing: (after quickly checking GPT) Yep. I just asked GPT to argue "Why AI will benefit the world" and "Why AI will destroy the world." Both answers sounded perfectly reasonable.

Muyao: That's the real problem we're facing now. If GPT can express both sides of a debate — and do it better than you — then we need to ask ourselves: is that debate even meaningful?

Ruo Han: Okay, before we end, I want to ask — on behalf of the listeners — in a world like this, how should we even live? How do we make peace with the modern life, and find a way to get through the rest of our days?

Muyao: There's this English phrase — "touch the grass." It means go live an offline life. In Chinese internet terms: breathe real air, pet real animals, eat what you truly love, hug the people you really want to hug.

Ruo Han: So basically… one day at a time? That's your advice, Muyao? Just live when it's time to live.

Muyao: One day at a time. Watch it all unfold like a play with no ending yet written. The world flipping upside down — it's kind of fascinating.

Wang Qing: Yeah. Let's leave it at that.

Ruo Han: Haha, beautiful. I'm off to soak up the Rio sunshine, hug the waves, eat the best Brazilian food I can find — and live in the moment, just like Teacher Muyao says.

Wang Qing: Exactly. That's it for today's episode. Hopefully we'll get to talk to Muyao again in a year, or two, or three — if we're still around.